

Developing A Cross-Cultural Annotation System and MetaCorpus for Studying Infants' Real
World Language Experience

Melanie Soderstrom (corresponding author)^a, Marisa Casillas^b, Erika Bergelson^c, Celia
Rosemberg^d, Florencia Alam^d, Anne S. Warlaumont^e, & John Bunce^f

^aDepartment of Psychology, University of Manitoba, Canada

^bMax Planck Institute for Psycholinguistics, Netherlands

^cDepartment of Psychology and Neuroscience, Duke University, USA

^dCentro Interdisciplinario de Investigaciones en Psicología Matemática y Experimental,
CONICET, Argentina

^eDepartment of Communication, University of California, Los Angeles, USA

^fDepartment of Human Development and Women's Studies, California State University, East
Bay, USA

Abstract

Recent issues around reproducibility, best practices, and cultural bias impact naturalistic observational approaches as much as experimental approaches, but there has been less focus on this area. Here, we present a new approach that leverages cross-laboratory collaborative, interdisciplinary efforts to examine important psychological questions. We illustrate this approach with a particular project that examines similarities and differences in children's early experiences with language. This project develops a comprehensive start-to-finish analysis pipeline that is publicly available for use, sensitive to cultural differences, and flexible to address a variety of research questions.

Key words: longform audio recordings, naturalistic observational research, interdisciplinary research

Introduction

The last 5-10 years have seen a massive rethinking of what it means to conduct research in psychology. Questions about the reproducibility of scientific findings (e.g. Klein et al., 2014), concerns about inherent cultural bias and the generalizability of our knowledge (e.g. Henrich, Heine, & Norenzayan, 2010), and the emergence of “Big Data” have created a perfect storm of change that continues to play out in different ways across sub-disciplines. While much attention has focused recently on experimental approaches (e.g. ManyLabs, Klein et al.; The ManyBabies Consortium, 2020), approaches that rely on naturalistic observations are equally impacted. Here we present an innovative, collaborative approach that leverages new technologies, naturalistic corpora and a conceptual coding system sensitive to diversity to begin to address basic questions about similarities and differences in human experience. Throughout the process, our focus was building a start-to-finish analysis pipeline, beginning from raw recorded observations, that would maximize the usefulness of our output beyond our specific project. We describe our project here with two goals: First, our novel annotation system and datasets will be of interest to other child language researchers, and potentially researchers from other fields. Second, we hope to inspire other researchers to apply our collaborative and interdisciplinary approach (as opposed to current mainstream siloed approaches) to other research questions.

The project described below coalesced around two fundamental challenges facing the community of researchers studying infants’ and toddlers’ language experiences. The first is one well-known to the broader psychological research community: research is heavily skewed toward certain specific populations that are not representative of the world as a whole, nor for that matter of human history (Henrich et al., 2010). Moreover, typical research approaches and

theoretical constructs are shaped by a narrow perspective about what is considered important in the mainstream culture (Bennis & Medin, 2010). This bias creates a narrow view of “typical” infant experience that problematically limits our understanding of the human condition. In language and language development research, this bias is compounded by basing assumptions for cross-cultural comparisons on one particular language: English (Majid & Levinson, 2010). This leads to conclusions driven by the unique properties of a single language and culture which are not representative of the diversity of human experience.

The second challenge has been a fundamental change in the basic building blocks of research in the sub-discipline of child language research. Until recently, methods and materials focused on painstaking hand-analysis of relatively short recordings of one-on-one interactions between caregiver and child. New technologies have emerged that allow researchers to record entire days of children’s typical lived experiences. The LENA (Greenwood, Thiemann-Bourque, Walker, Buzhardt, & Gilkerson, 2011) system is a pioneer in this area, providing lightweight, durable hardware and a software package that calculates estimates of basic measures of the child’s auditory language experience, such as the number of words spoken by adults throughout the day. This technological innovation has the potential to revolutionize our understanding of early language experience in ways that can help reduce cultural bias. To understand why, we must delve briefly into a debate within the language development community about the nature of early child language experiences.

A key tenet of mainstream child language research is that speech directed specifically at the child (referred to as “child-directed speech”, or CDS) holds a privileged role in the acquisition of language. The classical view of CDS holds as follows: Adults speak differently

when conversing with infants and young children than when speaking to other adults (e.g. Soderstrom, 2007). The particular linguistic and paralinguistic properties of CDS are beneficial for learning language by drawing attention to the speech signal (e.g. high/variable pitch), communicating affect (particular intonational structures, “happy talk”), and structuring the linguistic input in a form best-suited for learning (e.g. shorter utterances, simplified vocabulary). There is a large and varied literature supporting this view, and robust evidence that a) CDS exists in a broad spectrum of languages (e.g. Fernald et al., 1989), b) infants prefer CDS to other types of speech (The ManyBabies Consortium, 2020) and c) the quantity and quality of CDS is correlated with language outcomes (e.g. Rowe, 2012).

This “classical view” is straightforward, logical, and robustly supported by decades of research. However, while the literature presents a broad spectrum of languages where CDS has been documented, the extent to which children hear CDS varies radically across cultures (e.g. Brown & Gaskins, 2014; Cristia, Dupoux, Gurven, & Stieglitz, 2019; Gaskins, 2006; de León, 2011; Ochs & Schieffelin, 1984). Importantly, children around the globe successfully learn their language(s) despite this broad variation in CDS, unless they face clinical conditions that include language delays and disorders. However, past work estimating the quantity and properties of CDS children hear from caregivers has rarely taken a systematic cross-cultural/cross-linguistic approach (see, e.g., Broesch & Bryant, 2015; Fernald et al., 1989; Shneidman & Goldin-Meadow, 2012 for exceptions), particularly with respect to non-Western and small-scale traditional communities.

Direct cross-cultural or cross-linguistic comparisons of the process of linguistic development are fraught with issues of interpretability, in large part due to methodological

limitations. Until recently, researchers have been constrained in scope to what could be transcribed or annotated by hand. It is at best questionable, however, whether a one-hour recording with just mother and child in the home (a typical research approach) truly represents that child's real world language experiences. Cross-culturally, many infants are reported to seldom experience this one-on-one, single caregiver context (Brown & Gaskins, 2014; Gaskins, 2006; de León, 2011; Ochs & Schieffelin, 1984; Rosemberg, Alam, & Stein, 2014; Sperry, Sperry & Miller, 2018), and even in North America, a single hour of intensive interaction differs in important ways from a child's full day experience (Bergelson, Amatuni, Dailey, Koorathota, & Tor, 2019).

The introduction of daylong audio recordings (VanDam et al., 2016) as a methodological approach allows the researcher a much wider window into the child's experience, which can be highly variable over the course of a day (Anderson & Fausey, 2019). The LENA system (Greenwood et al., 2011) is a pioneer in this area, providing lightweight, durable hardware and a software package that calculates estimates of basic measures of the child's auditory language experience, such as the number of words spoken by adults throughout the day and child vocal assessment applied relatively quickly over multi-hour recordings. However, there remain hard limits that must be resolved before the promise of daylong recordings can be realized, particularly in cross-cultural work. For example, LENA does not differentiate child-directed from adult-directed speech, and may show variable performance across languages (Cristia, Bulgarelli, & Bergelson, 2019; Canault, Le Normand, Foudil, Loundon, & Thai-Van, 2016; Gilkerson et al., 2015; though see Cristia et al., 2020), having been developed only on North American English. This leads not only to concerns about accuracy, but also to the possibility of

introducing systematic bias in a comparative analysis. Moreover, LENA is proprietary software that to the best of our knowledge is not currently undergoing significant further development to improve its automatic labeling algorithms. A multipurpose speech processing tool, one that is applicable across languages and cultural contexts, must be built. Such a tool cannot emerge *de novo*, but must be developed via machine learning over a representative set of hand-tagged audio recordings on a scale not typical of individual research.

Automated analysis in itself cannot resolve cultural bias - it will simply reflect the biases built into the system, including biases in the structure of the training data. To conduct a comparative analysis that reduces bias requires not only building new tools, but developing a new framework for analysis (an annotation system, a sampling approach, etc.) that can be applied across a broad spectrum of lived experiences. What was needed was a pipeline for taking longform audio recordings of children's real everyday language experiences and creating a comprehensive dataset available to answer a variety of questions about those experiences, as well as providing the primary data source to leverage the creation of automated tools that could apply these same questions across a much larger dataset.

ACLEW (Analyzing Child Language Experiences around the World) was therefore conceived by a group of child language ("datasets") researchers who sought to partner with a group of speech technology ("tools") researchers, working collaboratively to address these needs. Working in our favour, the child language research community has a long history of collaborative data-sharing (MacWhinney, 2000; Vandam et al., 2016), that has been adopted more recently in the broader developmental community (Simon, Gordon, Steiger, & Gilmore, 2015). Creating a cross-cultural and cross-linguistic dataset requires researchers to share their

data (“corpora”) across the traditional laboratory siloes. Each datasets investigator contributed a set of recordings originally collected in their own laboratory for their own specific research question, to create a larger, cross-linguistic and cross-cultural meta-corpus. With this combined dataset in hand, the following objectives were pursued:

Goal 1: Build a flexible but systematic annotation system (ACLEW Annotation Scheme)

Existing transcription systems were designed for short recordings and to examine specific linguistic details within a language, not for comparing language experience in daylong recordings across language communities. It was therefore necessary to create an annotation system that could be applied to a large number of research questions, but that would provide enough standardization to ensure that the individual corpora were being measured similarly despite differences in culture, language and sampling across the corpora.

Goal 2: Implement sampling across language communities (ACLEW dataset)

Once the annotation was developed, a representative sample across the corpora was extracted based on a number of desiderata. The annotation system was then implemented on this sample to create the core dataset for analysis. Our hope is that this dataset will be the starting point for future contributions using this framework.

Goal 3: Provide a well-designed corpus for the development of tools for automated analysis

While the hand-annotated dataset can be used to directly address questions about the language experiences of infants and young children, an important objective of this project was to support

the development of tools to analyze the samples in a more automated fashion. Automated tools allow researchers to gain a wider perspective than can be gained by a hand-sampling approach, but pose a huge technical challenge to develop even with state-of-the-art speech processing technologies.

The objective was not simply to develop procedures and standards for a single self-contained project, but to provide a streamlined, start-to-finish pipeline that can be implemented easily and widely in the same way across many and diverse laboratories. This approach can leverage the many thousands of hours of such audio recordings being collected in laboratories around the world (for example, HomeBank, a repository of daylong child-centered audio recordings, currently contains approximately 12,000 hours of audio data, and this is only a small subset of recording hours potentially available for use) to begin to address questions that require a larger and more diverse dataset than could be obtained via one or a small number of laboratories, using a shared set of measures that can help answer a broad variety of specific research questions. An early attempt at such a collaborative approach (Bergelson, Casillas, Soderstrom, Seidl, Warlaumont, & Amatuni, 2019), was limited to North American samples and a very barebones classification system, but showed the promise of this collaborative framework. More recently, a much broader-based analysis (Bergelson et al., 2020) leveraged many thousands of hours of LENA recordings from diverse communities to examine the variation in exposure to adult speech and its relationship to the quantity of infant vocalizations.

In the remainder of the paper, we describe the development of the annotation scheme and the ACLEW dataset in detail. Along the way we articulate a number of questions and challenges

that arose, with a focus on those that would be of most broad interest. Finally, we summarize the benefits and limitations of this approach, and discuss the implications and insights for other research programs. In addition to introducing the ACLEW dataset and annotation scheme as resources in and of themselves, we hope to inspire other researchers, both within and outside of the child language research community, to join in our larger objectives of building a comprehensive, diverse and detailed collection of source data for the many theoretical questions and needed tool developments that remain.

The ACLEW Annotation Scheme

In designing the annotation system, there were a number of complementary and competing objectives. The system needed to be easy and relatively quick to implement across a large number of laboratories on longform audio recordings. A comprehensive set of tutorials were created for all components of the system (links below), as well as a system for measuring new trainees' performance against a set of "gold standard" (ideal) practice files. A high degree of structure in the design was critical, both for comparability across laboratories, but also because the annotated files are intended to be machine-readable. In addition, however, it was important that the system be adaptable to a variety of research questions so that other researchers outside of the formal ACLEW collaborative (or even outside the child language research community) could use the system and in some cases share their data with tools developers. Indeed, some laboratories outside of ACLEW have already begun to use versions of this system. It was also important that the system be compatible with existing systems such as the widely-used CHAT transcription system (MacWhinney, 2000). Luckily, there is already a large degree of

interoperability across existing transcription/annotation platforms. Indeed, the use of the ELAN system (Sloetjes & Wittenburg, 2008) as our base permits easy export to CLAN (MacWhinney, 2000, a tool for analysis of CHAT transcripts) or Praat (Boersma, 2001, a tool for acoustic/phonetic analysis). Lastly, it was important that the system itself minimize cultural bias in the design.

This latter desideratum proved challenging in some ways that were especially informative of our larger research goals, both from the perspective of cross-cultural comparison and in terms of technical considerations. For example, it is common practice in analyses of North American samples to exclude “naptimes”, when the infant is sleeping, if the recorder is left on. However, “naptime” as a construct does not exist universally across cultures. For example, Mayan infants sleep at times of their own choosing (Morelli, Rogoff, Oppenheim, & Goldsmith, 1992), and are often carried (Brown, 2014) in a soothing atmosphere during which dozing may occur (De León, 2011; Pye, 1986), making naptime at best hard to identify within a recording, and arguably culturally inappropriate. Moreover, failure to identify and transcribe speech present in the recording while the target child was sleeping would be troublesome for the speech detection algorithms.

Treatment of television within the recordings similarly created difficulties due to systematic variation in its presence across corpora, and the technical limitations of speech processing to distinguish live from recorded speech. Other differences across the corpora that led to challenges in decision-making (but in ways meaningful to our larger research questions) included the number of speakers present at any given time and the locations (e.g. indoors, outdoors) in which activities took place. The project also forced structured examination of key

theoretical constructs that often vary across projects, such as defining “child-directed speech”, and considering metrics of children’s own productions that would be neutral to cross-linguistic differences (e.g. a given language’s phonetic inventory).

Components of the ACLEW Annotation Scheme

System Overview.

The annotation scheme (<https://osf.io/aknjv/>) is built on a simpler framework known as the DARCLE Annotation Scheme (DAS; Casillas et al., 2017; <http://osf.io/4532e>).¹ The DAS provides a minimal framework for dividing the audio stream into labeled and unlabeled segments across a set of speaker tiers using ELAN, a media annotation application (Sloetjes & Wittenburg, 2008). The DAS is distributed as an ELAN template file that can be adapted for individual recordings and research projects. In the ACLEW adaptation of the DAS templates (the ACLEW DAS, or AAS), further template structure is added to annotate more information about the labeled speaker segments in a series of subtiers (Figure 1).

In brief, we first tag who’s speaking when. For the vocalizations of the “target child” (the child whose language environment is under study, who is wearing or close to the recorder), we add further classifications. Depending on the age of the child, we add metrics of vocal maturity that are cross-linguistically applicable (e.g. whether vocalizations contain canonical or non-canonical syllables; see Figure 1). For all speakers other than the target child, the intended addressee is also indicated in a subtier (adult, child, both or unsure). Finally, speech is

¹ DARCLE is an open consortium on **Daylong Audio Recordings of Children’s Language Environments**; see darcle.org

transcribed using the minCHAT format, used in collections of child language recordings going back to the 1980's (MacWhinney, 2000).

CHI	ba [: ball] Mama.	0.	0.
—vcm@CHI	C	N	C
—lex@CHI	W		0
—mwu@CHI	M		
FA1		wait wait it's right here.	
—xds@FA1		C (T)	
MA1	hey-.		
—xds@MA1	A		
MC1	no no no no.	<xxx> [=! sings].	
—xds@MC1	U	U	

Figure 1. Example annotations for a few seconds of audio. In this fictional clip, a female adult (FA1) speaks directly to the target child (CHI), a male adult (MA1) begins to speak to the female adult but cuts himself off, and a child is talking and singing, but it is unclear who his addressee is. All CHI utterances are annotated for vocal maturity (i.e. given canonical babble, check for use of recognizable words; given recognizable words, check for multi-word use). '0.' is used as a placeholder for untranscribed utterances. All non-CHI utterances are annotated for addressee type (C = child-directed speech / T can be used instead to indicate target-child directed speech specifically).

Tutorials.

The AAS is described in detail in a series of self-paced training tutorials (<https://osf.io/b2jep>), available in both English and Spanish, for Mac and PC, providing template files, details and step-by-step instructions on implementing the various components of the system, and a final process for checking annotation files. Lastly, we have a series of gold standard training and test files in both English and Spanish, against which new trainees are compared. Each Gold Standard file was initially annotated and transcribed separately by two experienced annotators who then created a consensus file. This file was reviewed by two supervisory annotators to create a final consensus Gold Standard file. This was a crucial part of our process to ensure comparability of the annotations of research assistants across laboratories. This is a uniquely comprehensive, publicly available training system for annotation of audio files, including a web app to automate comparison to the Gold Standard (<https://github.com/aclew/ACLEW-GoldStandard>). It compares segmentation overlap and accuracy of different tags and provides a weighted score indicating whether the trainee has passed, as well as an error report indicating specific discrepancies.

ACLEW Dataset

The ACLEW Dataset is a meta-corpus of 7 individual corpora from different research laboratories, that were each collected for different research questions related to child language development (Bergelson, 2016; Casillas et al., 2019; McDivitt & Soderstrom, 2017; Rowland, Bidgood, Durrant, Peter, & Pine, 2018; Rosemberg, Alam, Stein, Migdalek, Menti, & Ojea, 2015; Vandam et al., 2016; Warlaumont, Pretzer, Mendoza, & Walle, 2016). All data were collected with the approval of each laboratory's supervising ethics board. Table 1 provides demographic and other information for each corpus. Importantly, while this collection is not

representative of the diversity worldwide of human experiences and languages, it does provide a snapshot with sufficient variation to push the limits of an annotation system. The corpora include both industrialized and small scale societies, diverse socio-economic situations, and linguistically unrelated languages.

Table 1. Details of the full corpora available in the ACLEW dataset. Some files (McDivitt, Warlaumont and some Bergelson and Casillas files) are available through the HomeBank datasharing system. The remaining files may be requested directly from the dataset owner.

	Language(s)	Ages	Region	Participants Recordings Hours	Recording Devices
Bergelson	American English	6–17mo	NE U.S.	44 522 >7300	LENA™
Casillas- Tseltal	Tseltal (Mayan)	2–36 mo	Chiapas, Mexico	10 10 >90	Olympus WS-832
Casillas- Yéli	Yéli Dnye (Papuan)	1–36 mo	Milne Bay Province, Papua New Guinea	10 10 >80	Olympus WS-832 and Olympus WS-835
Lang0-5	UK English	11–31 mo	Great Britain	39 235 >3500	LENA™
McDivitt/ Winnipeg	Canadian English and Canadian French	2–34 mo	Winnipeg, Manitoba, Canada	12 43 >350	LENA™
Rosemberg	Argentinian Spanish	8–35 mo	Buenos Aires,	20 60	Different digital

			marginal slums, Argentina	>240	devices: Sony, Panasonic, RCA, Olympus, LENA™
Warlaumont	American English and Spanish	3–18 mo	Central Valley California, US	15 40 >490	LENA™
TOTAL		1–36 mo		150 920 >12050	

Data sharing.

One of the biggest challenges to resolve in a project of this type is balancing a commitment to Open Science with legitimate concerns across the corpora regarding issues of participant confidentiality and consent to share. The corpora contain recordings of the intimate details of participants’ daily lives, and potentially unconsented third parties. Essentially by definition (given current and future voice recognition technology, etc.) it is not possible to fully “anonymize” such daylong recordings. In addition, the corpora include at-risk and vulnerable individuals, and communities that could be significantly impacted as a group by negative evaluation or misunderstanding. Furthermore, because the project relied on pre-existing data sources, there was diversity in the type and extent of participant consent for sharing.

During the course of the project, access to the full corpora has been therefore limited to specific members of each laboratory as needed to conduct their research. All researchers with access to the dataset undertook an ethics tutorial, and the Principal Investigator of each

laboratory signed a memorandum of understanding regarding data usage and sharing (<https://osf.io/erkm8/>). While tools were agreed to be fully public as risk to participants is minimal, the raw audio recordings and the annotations, transcriptions, metadata, and derivative data from later stages in the analysis pipeline were kept private to varying extents (as summarized in Table 2).

Table 2. Data sharing ground rules

	Data type	Access during ACLEW project
Raw data	Audio files	Private
Data codified	Transcriptions	Private
	Annotations	Private
Metadata	Individual	Private
	Group level characteristics (<i>e.g. descriptive features, e.g. group level data on gender, education level</i>)	Public
Derivative data	Fully anonymized data from the audio and/or annotation files that feeds into a publication pipeline	Public
	Anonymized data at the individual level, feeding into a publication pipeline, but that involves a marginalized group and a potential ethical concern is articulated	Restricted

These rules outlined in the memorandum defined a minimal level of data sharing during the time of the project, depending on the needs of each individual investigator responsible for a dataset; individual investigators could provide less restrictive access at their discretion. It is crucial for projects of this type to carefully consider these concerns from the beginning and outline explicitly the expectations for data sharing. There is a growing set of resources available to help navigate these concerns (e.g. Casillas & Cristia, 2019; Cychosz et al. 2020).

A major practical concern was what software platform was best for enabling effective data sharing. We focus here on *how* the decision was reached, rather than the platform(s) ultimately used, given the rapid rate at which platforms for sharing data are currently evolving. It was necessary for the system to provide a high level of security for confidential files, but also give flexible and easy access by a relatively large and changing set of researchers and assistants. The dataset includes both static, longform data (i.e. 2 GB individual daylong audio recordings), but also derivative annotation files (manual and algorithm-derived) for which version control is a concern. Easy upload and download in batch form is important, and this was particularly challenging to resolve due to the size of the dataset (about 150 GB of annotated files, and over 1 TB of raw, unannotated audio). Audio recordings exceed the current limit for private storage on GitHub (<https://github.com>). On other platforms, space saving processes (conversion to mp3 from wav) generated concerns about maintaining original recording quality. It is also necessary to consider changing needs for access both during and after the project. Table 3 provides a summary of the key considerations.

Table 3. Checklist for data-sharing and storage.

Security	Adequate password protection for confidential/sensitive files.
Access	Access must be flexible but controlled to accommodate changing research assistants and researchers within a large collaborative framework.
Version control	Complex workflow requires careful consideration of version control for derived data and annotations.
Size	Must be able to accommodate very large datasets.
Fidelity	Storage must be in original format.
Flexibility	Types of access, etc. must allow for changing needs over the course of the project.

Sampling and Implementation.

Practicality dictates that only a subset of the files can be annotated by hand (hence the need to develop automated tools). However, the act of selecting subsections inherently narrows the scope of analysis, and introduces potential for bias. Selection occurs at multiple levels: participants within a corpus, recordings from a given participant, and time windows within a recording (time of day and length of selection, whether a context window is provided around the selection, etc.).

Decisions at these various levels may be made for practical reasons (e.g. limitations on available work hours, sharing permissions and sensitive content), to equate characteristics across corpora or participants (or longitudinally within a participant's data), or to ensure diversity across a dataset. In all cases, these decisions will have important consequences for the nature of the conclusions that can be drawn, particularly when comparing across datasets that are heterogeneous in many of the relevant characteristics (length of recordings, number of recordings, longitudinal or cross-sectional data collection, etc.), as was the case for ACLEW.

An initial pilot sample, *starter-ACLEW* (<http://doi.org/10.17910/B7.390>), was focused on selecting short audio clips to pilot and refine the annotation system. Three 5-minute samples were hand-selected from each corpus (except Casillas-Yéli, which was not yet available), with the goal of obtaining clips that were reasonably representative of the recordings within the corpora, but with sufficient quantity and diversity of adult and child vocalizations and diversity of child age.

A first round of annotation was then implemented by selecting up to 10 recordings from each corpus to reflect its maximal spread in child sex, maternal education and child age (0–36 months). A *fully random sampling* approach was taken in each recording, in which 15 short 2-minute samples within a 5-minute context window were randomly selected by algorithm within each selected file (except for the Casillas files, where 9 5-minute (Tseltal) or 2.5-minute (Yéli samples were selected, see below, Table 4). This approach was chosen rather than, for example, repeated sampling at regular intervals due to the diverse nature of the recordings, which varied with respect to recording length and time-of-day.

However, while random sampling minimized the potential impact of some kinds of analytic bias, it introduced other concerns, most notably that the data were highly zero-inflated - in many cases the sampled selection contained no speech to annotate. While this was important information from the perspective of understanding infants' holistic language environment, it created difficulties for analysis of the speech content. This skewed distribution was problematic in selecting an appropriate statistical approach for comparison, and also because the total sample of each speech type was relatively low, which impeded comparisons across classification categories (e.g., speech from male vs. female adults). A second round of annotation was therefore implemented, using **high-volubility sampling**. In this approach, samples were preselected for high speech volubility using an early version of the automated tool developed by our machine learning collaborators, DiViMe (Le Franc et al., 2018; see below). DiViMe generated candidate audio clips in which it estimated that a lot of speech was occurring. Human listeners then screened these clips to exclude instances where the tool mis-identified non-speech as speech (e.g. television, crying, heartbeats; <https://osf.io/739g8/wiki/home/>).²

Table 4. Descriptive data from random and high volubility sampling of the ACLEW dataset. *R* = random sample, *H* = high volubility sample. *Total min. transcribed* refers to the total minutes of audio sampled for transcription. *Total min. of speech* is the total minutes of audio identified as containing speech. *ADS* is the quantity of speech produced that was adult-directed (in seconds per minute to standardize across corpora with different sampling quantities). *CDS* is the quantity

² Note that the Casillas Tseltal and Yélî corpora follow a different approach for sampling (see Casillas, Brown, & Levinson, 2020) due to the practical limitations of working locally with Indigenous informants on clip annotation/transcription in these two language communities.

of speech produced that was infant-directed. Note that total speech (but not ADS/CDS) includes target child vocalizations and avoids double-counting overlapping speech segments. Speech from electronic media sources is annotated in the source data but not included in these tallies.

		Total min. of speech / Total min. sampled	ADS (s/min)	CDS (s/min)
Bergelson	R	83.4 / 300	8.54	5.15
	H	137.73 / 300	10.37	10.26
Casillas- Tseltal	R	189.64 / 450	11.76	8.76
	H	61.51 / 150	4.81	10.70
Casillas- Yéli	R	139.65 / 225	21.08	17.45
	H	87.05 / 150	15.83	15.92
Lang0-5	R	77.86 / 300	4.34	5.42
	H	143.27 / 300	4.09	14.22
McDivitt/ Winnipeg	R	82.32 / 300	6.09	4.98
	H	67.45 / 180	5.32	10.27
Rosemberg	R	122.51 / 300	10.78	7.21
	H	98.54 / 270	7.10	8.06
Warlaumont	R	81.37 / 300	8.28	4.14
	H	95.06 / 210	9.10	11.34

Reliability.

A reliability check was performed across laboratories on the random sample to ensure cross-laboratory consistency. One minute from one of the annotated 2-minute segments in each file was selected, omitting any segments with no speech. Further details are available on osf (<https://osf.io/vbpqf>).

A sneak peek at tools for automated processing of longform audio

Given our present goal of high-level description of our collaborative approach centered on child language, we provide only a brief description of the automated speech processing components of the project, as an illustration of the importance of this kind of collaborative work. The interested reader is directed to recent publications from our “tools” colleagues (DiViMe: Le Franc et al., 2018; ALICE: Räsänen et al., 2019, 2020, see also [https://divime.readthedocs.io/en/latest/³\)](https://divime.readthedocs.io/en/latest/³)) for more detailed information. In short, progress has been made in providing tools that answer the following basic annotation questions: (a) where in the recording is speech/vocalization happening? (b) what type of speaker is speaking (e.g., a man, a child, etc.)?, (c) what level of vocal maturity is this child vocalization (e.g., noncanonical babble? canonical babble)?, (d) and how many syllables/words are there in this stretch of adult speech?

The ongoing collaborative relationship between child language and tools researchers factored heavily into the ACLEW work process. For example, for the sake of reducing workload on research assistants, the original plan called for full transcription only of speech directed at children, with only basic segmentation/annotation provided for adult-directed speech. Early work on implementation of the syllabifier made it clear, however, that transcription of adult-directed speech was necessary for accurate tool development. Such collaborations can be challenging given different norms across disciplines, but are critical for progress.

³ Note: the DiViMe system is no longer being actively supported, but works with some operating systems. We anticipate a newer system to emerge in the near future based on the latest developments with ALICE.

Concluding Thoughts

The preceding pages provide an overview both of the process and the product in the ACLEW project. We hope that both will be of use to researchers both within and outside of child language research as we turn toward naturalistic observational work that spans full days of human experience. This project has been a unique experience for all involved, wide-ranging in its efforts to build a framework for cross-cultural comparative analysis, pushing forward new methods and new technologies, and planting the seeds for more robust cross-disciplinary research.

It is important to acknowledge that the dataset to date, while taking a few steps beyond the typical North American-centric approach, remains limited to a small number of languages and cultural contexts (not to mention researcher backgrounds). Furthermore, it focused only on spoken language, which excludes both gestural and gaze information within a spoken language context, and language in other modalities. While acknowledging that these are only modest first efforts, these corpora represent a highly diverse sample, and constitute a unified dataset that allows for more meaningful, apples-to-apples, comparisons across language communities that will be indispensable in making progress in understanding the true range of circumstances under which human children learn their community's language(s). Moreover, and importantly, this work lays the foundation on which others can build, on this and other research topics.

We encourage other researchers to use, and build on, the ACLEW Annotation Scheme framework. Not only does it come with self-guided tutorials and a ‘test’ for research assistants to add quality control and standardization, it provides the potential for more direct comparisons across diverse communities, makes data interoperable and extendible beyond their original collection purposes, and increases the potential usefulness of all researcher’s audio recordings

for ongoing and future development of automated speech processing tools. Notably, our machine learning colleagues report that one of the biggest impediments to progress is the lack of a sizeable corpus of consistently-tagged and carefully segmented audio recordings of children's real-world experiences. A small commitment of effort from the child language research community would leverage significant benefits from theirs.

Over the course of the project, a number of technical and decision-making challenges emerged that are inherent in this approach. Modern technology provides many tools for collaboration across geographical distance, but does not address the considerable diversity across laboratories in experiences, perspectives, objectives, work processes and resources. These differences led to different collaborative approaches to important issues such as data sharing, workflow, and timelines, which often led to a considerable investment of time in discussion and careful documentation, but ultimately also to valuable insights. For example, discussions around cultural practices such as naptime and television usage, and best approaches to measuring child vocalization, were in and of themselves of theoretical interest. Our discussions also led us to tackle challenging ethical concerns in a meaningful way - for example our data sharing policy is a careful compromise across laboratories with very different perspectives on data sharing.

The cross-disciplinary approach was also of critical value not only in terms of the promise of better automated tools down the road. Across many joint meetings, we were able to clarify the needs and priorities of researchers across what has traditionally been a fairly wide divide. With some notable exceptions (e.g. Fell, Cress, MacAuslan, & Ferrier, 2004; Oller et al., 2010; Ramsay, Ghai, Kumareswaran, Edwards & Bailey, 2019; Warlaumont & Ramsdell-Hudock, 2016), those interested in using machine learning to develop automated

speech processing tools for real-world audio, and the child language researchers generating the underlying datasets over which the machine learning is run, have not been in direct communication. Our ongoing discussions have allowed child language researchers to better understand what makes a “good” dataset from the perspective of tool-building, and machine learning experts to understand what kinds of tools may benefit the child language research community.

The ACLEW annotation scheme, dataset, and speech tools will allow us to address many substantive questions about the quantitative and qualitative nature of children’s real world experiences and ultimately to generate insights that will inform our understanding of the diversity (and similarity) of children’s early language experiences. The goal of this project was to provide a stepping stone for further inquiry, both in broadening the scope of language communities and in providing a framework for addressing other important questions, such as the nature of caregiver-child interactions, expression of vocal affect, experiences with media, etc. Many of these questions could be readily examined within the existing AAS - others may require expansion and adaptation of the current framework. We provide a number of publicly available tools that can be leveraged for a variety of extensions to our approach. Importantly, we also illustrate a model for cross-laboratory, cross-cultural, international and cross-disciplinary research that we hope will be of interest to researchers in other sub-disciplines.

References

- Anderson, H.T., & Fausey, C. (2019, March). *Modeling non-uniformities in infants' everyday speech environments*. Paper presented at Society for Research in Child Development 2019 Biennial Meeting, Baltimore, MD.
- Bergelson, E. (2016). Bergelson seedlings homebank corpus. doi:10.21415/T5PK6D.
- Bergelson, E., Amatuni, A., Dailey, S., Koorathota, S., & Tor, S. (2019). Day by day, hour by hour: Naturalistic language input to infants. *Developmental science*, 22(1), e12715.
- Bergelson, E., Casillas, M., Soderstrom, M., Seidl, A., Warlaumont, A. S., & Amatuni, A. (2019). What do North American babies hear? A large-scale cross-corpus analysis. *Developmental science*, 22(1), e12724.
- Bergelson, E., Soderstrom, M., Schwarz, I.-C., Rowland, C., Ramirez-Esparza, N., Rague Hamrick, L., Marklund, E., Kalashnikova, M., Guez, A., Casillas, M., Benetti, L., van Alphen, P., & Cristia, A. (2020, submitted). Everyday language input and production in 1000 children from 6 continents.
- Bennis, W. M., & Medin, D. L. (2010). Weirdness is in the eye of the beholder. *Behavioral and Brain Sciences*, 33(2-3), 85-86.
- Boersma, Paul (2001). Praat, a system for doing phonetics by computer. *Glott International* 5(9/10), 341-345.
- Broesch, T. L., & Bryant, G. A. (2015). Prosody in infant-directed speech is similar across Western and Traditional Cultures. *Journal of Cognition and Development*, 16(1), 31-43.

- Brown, P. (2014). The interactional context of language learning in Tzeltal. In I. Arnon, M. Casillas, C. Kurumada, & B. Estigarriba (Eds.), *Language in Interaction: Studies in honor of Eve V. Clark* (pp. 51-82). Amsterdam: Benjamins.
- Brown, P., & Gaskins, S. (2014). Language acquisition and language socialization. In *Cambridge handbook of linguistic anthropology* (pp. 187-226). Cambridge University Press.
- Canault, M., Le Normand, M. T., Foudil, S., Loundon, N., & Thai-Van, H. (2016). Reliability of the language environment analysis system (LENA™) in European French. *Behavior Research Methods*, 48(3), 1109-1124.
- Casillas, M., Bergelson, E., Warlaumont, A. S., Cristia, A., Soderstrom, M., VanDam, M., & Sloetjes, H. (2017). A new workflow for semi-automatized annotations: Tests with long-form naturalistic recordings of children's language environments. *INTERSPEECH 2017*, 2098-2102.
- Casillas, M., Brown, P., & Levinson, S. C. (2019). Early language experience in a Tseltal Mayan village. *Child Development*.
- Casillas, M., & Cristia, A. (2019). A step-by-step guide to collecting and analyzing long-format speech environment (LFSE) recordings. *Collabra*, 5(1).
- Cristia, A., Bulgarelli, F., & Bergelson, E. (2020). Accuracy of the Language Environment Analysis System Segmentation and Metrics: A Systematic Review. *Journal of Speech, Language, and Hearing Research*, 63(4), 1093-1105.

- Cristia, A., Dupoux, E., Gurven, M., & Stieglitz, J. (2019). Child-directed speech is infrequent in a forager-farmer population: A time allocation study. *Child Development, 90*(3), 759-773.
- Cristia, A., Lavechin, M., Scaff, C., Soderstrom, M., Rowland, C., Räsänen, O., ... & Bergelson, E. (2020). A thorough evaluation of the Language Environment Analysis (LENA) system. *Behavior Research Methods*.
- Cychosz, M., Romeo, R., Soderstrom, M., Scaff, C., Ganek, H., Cristia, A., ... & Weisleder, A. (2020). Longform recordings of everyday life: Ethics for best practices. *Behavior Research Methods*.
- De León, L. (2011). Language socialization and multiparty participation frameworks. In: A. Duranti; E. Ochs & B.B. Schieffelin (Eds.) *The Handbook of Language Socialization* (pp. 81-111). United Kingdom: Wiley-Blackwell.
- Fell, H. J., Cress, C., MacAuslen, J., & Ferrier, L. J. (2004). visiBabble for reinforcement of early vocalization. *ASSETS*, 161-168.
- Fernald, A., Taeschner, T., Dunn, J., Papousek, M., de Boysson-Bardies, B., & Fukui, I. (1989). A cross-language study of prosodic modifications in mothers' and fathers' speech to preverbal infants. *Journal of child language, 16*(3), 477-501.
- Gaskins, S. (2006). Cultural perspectives on infant-caregiver interaction. In N. J. Enfield and S. C. Levinson (Eds.) *The roots of human sociality: Culture, cognition, and interaction* (pp. 279–298). Oxford, UK: Berg.

- Gilkerson, J., Zhang, Y., Xu, D., Richards, J. A., Xu, X., Jiang, F., ... & Topping, K. (2015). Evaluating Language Environment Analysis system performance for Chinese: A pilot study in Shanghai. *Journal of Speech, Language, and Hearing Research*, 58(2), 445-452.
- Greenwood, C. R., Thiemann-Bourque, K., Walker, D., Buzhardt, J., & Gilkerson, J. (2011). Assessing children's home language environments using automatic speech recognition technology. *Communication Disorders Quarterly*, 32(2), 83-92.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61-83.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., & Bernstein, M. J. ... & Nosek, B.A. (2014). Investigating variation in replicability: A “many labs” replication project. *Social Psychology*, 45(3), 142-152.
- Le Franc, A., Riebling, E., Karadayi, J., Wang, Y., Scaff, C., Metze, F., & Cristia, A. (2018). The ACLEW DiViMe: An easy-to-use diarization tool. In B. Yegnanarayana, C. Chandra Sekhar, S. Narayanan, S. Umesh, S. R. M. Prasanna, Hema A. P. Kumar Ghosh, et al. (Eds.) *Proceedings of Interspeech 2018* (pp. 1383–1397). DOI: <https://doi.org/10.21437/Interspeech.2018-2324>
- Majid, A., & Levinson, S. C. (2010). WEIRD languages have misled us, too. *Behavioral and Brain Sciences*, 33(2-3), 103-103.
- ManyBabies Consortium, The. (2020). Quantifying sources of variability in infancy research using the infant-directed speech preference. *Advances in Methods and Practices in Psychological Science*, 3(1), 24-52.

- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk. Third Edition.*
Mahwah, NJ: Lawrence Erlbaum Associates.
- McDivitt, K., & Soderstrom, M. (2016). McDivitt homebank corpus. doi:10.21415/T5KK6G.
- Morelli, G. A., Rogoff, B., Oppenheim, D., & Goldsmith, D. (1992). Cultural variation in infants' sleeping arrangements: Questions of independence. *Developmental Psychology*, 28(4), 604.
- Ochs, E., & Schieffelin, B. (1984). Language acquisition and socialization: Three developmental stories and their implications. *Culture theory: Essays on mind, self, and emotion*, 276-320.
- Oller, D. K., Niyogi, P., Gray, S., Richards, J. A., Gilkerson, J., Xu, D., Yapanel, U., & Warren, S. F. (2010). Automated vocal analysis of naturalistic recordings from children with autism, language delay, and typical development. *Proceedings of the National Academy of Sciences of the United States of America*, 107(30), 13354-13359.
- Pye, C. (1986). Quiché Mayan speech to children. *Journal of Child Language*, 13(1), 85-100.
- Ramsay, G., Ghai, S., Kumareswaran, M., Edwards, M., & Bailey, J. (2019). Developmental cascades in reciprocal vocal signaling between infant and caregiver in typical development and autism. *The Journal of the Acoustical Society of America*, 145(3), 1731-1731.
- Räsänen, O., Seshadri, S., Karadayi, J., Riebling, E., Bunce, J., Cristia, A., ... & Soderstrom, M. (2019). Automatic word count estimation from daylong child-centered recordings in various language environments using language-independent syllabification of speech. *Speech Communication*, 113, 63-80.

- Räsänen, O., Seshadri, S., Lavechin, M., Cristia, A., & Casillas, M. (2020). [ALICE: An open-source tool for automatic measurement of phoneme, syllable, and word counts from child-centered daylong recordings](#). *Behavior Research Methods*.
- Rosemberg, C.R., Alam, F. & Stein, A. (2014). Factors reflecting children's use of temporal terms as a function of social group. *Language, Interaction and Acquisition*, 5(1), 38–61.
- Rosemberg, C.R., Alam, F., Stein, A., Migdalek, M., Menti, A., & Ojea, G. (2015). Los entornos lingüísticos de niñas y niños pequeños argentinos / language environments of young argentinean children. CONICET (DOI in progress).
- Rowe, M. L. (2012). A longitudinal investigation of the role of quantity and quality of child-directed speech in vocabulary development. *Child development*, 83(5), 1762-1774.
- Rowland, C.F., Bidgood, A., Durrant, S., Peter, M. Pine, J.M. (2018). The Language 0-5 Project, 10.17605/OSF.IO/KAU5F Unpublished manuscript available from <https://osf.io/kau5f/>.
- Sloetjes, H., & Wittenburg, P. (2008). *Annotation by category – ELAN and ISO DCR*. In Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008).
- Soderstrom, M. (2007). Beyond babytalk: Re-evaluating the nature and content of speech input to preverbal infants. *Developmental Review*, 27(4), 501-532.
- Sperry, D. E., Sperry, L. L., & Miller, P. J. (2019). Reexamining the verbal environments of children from different socioeconomic backgrounds. *Child Development*, 90(4), 1303-1318.

- Simon, D. A., Gordon, A. S., Steiger, L., & Gilmore, R. O. (2015, June). *Databrary: Enabling sharing and reuse of research video*. In Proceedings of the 15th ACM/IEEE-CS Joint Conference on Digital Libraries (pp. 279-280). ACM.
- VanDam, M., Warlaumont, A. S., Bergelson, E., Cristia, A., Soderstrom, M., De Palma, P., & MacWhinney, B. (2016). HomeBank: An online repository of daylong child-centered audio recordings. *Seminars in speech and language, 37*(2), 128-142.
- Warlaumont, A.S., Pretzer, G.M., Mendoza, S. & Walle, E.A. (2016). Warlaumont homebank corpus. doi:10.21415/T54S3C
- Warlaumont, A. S., & Ramsdell-Hudock, H. L. (2016). Detection of total syllables and canonical syllables in infant vocalizations. *INTERSPEECH 2016*, 2676-2680.

Author Contributions

Contributed to conception and design: MS, MC, EB, CR, FA, AW, JB

Contributed to acquisition of data (audio and/or annotation): MS, MC, EB, CR, FA, AW, JB

Contributed to analysis and interpretation of data: N/A

Drafted and/or substantially revised the article: MS, MC, JB

Approved the submitted version for publication: MS, MC, EB, CR, FA, AW, JB

Acknowledgments

The authors wish to thank Caroline Rowland and the creators of the Lang0-5 corpus for generously providing their time and access to their recordings, and to our “tools” colleagues: Alex Cristia, Okko Räsänen, Florian Metze, and Björn Schuller and their teams. Thanks are also due to the many participants who agreed to share their recorded experiences, and to the volunteer and paid research assistants across several laboratories who contributed to data collection and to the development and implementation of the ACLEW Annotation Scheme.

Funding

This work was supported by a Trans-Atlantic Platform “Digging into Data” collaborative grant (ACLEW: Analyzing Child Language Experiences Around The World; HJ- 253479-17 (EB), SSHRC-869-2016-0003(MS), NSERC-501769-2016-RGPDD(MS), MINCyTHJ-253479(CR)). We acknowledge further support from Social Sciences and Humanities Research Council of Canada Insight Grant (435-2015-0628; MS); National Institutes of Health (DP5 OD019812-01; EB), CONICET grant PIP 80/201 and SECyT grant PICT 3327/2014 (CR), National Science

Foundation grants 1529127 and 1827744 (formerly 1539129) and the James S. McDonnell Foundation (ASW).

Competing Interests

The author(s) declare that there were no competing interests with respect to the authorship or the publication of this article.

Data Accessibility

Detailed methodological information, tutorials and analysis scripts can be found on our project page as listed in the manuscript. ACLEW scripts can be found at repositories on github: <https://github.com/aclew>. Individual datasets can be accessed either through HomeBank (<https://homebank.talkbank.org/>) or via permission from the dataset owner. Due to privacy concerns for the participants, full audio is not publicly available.

Figures

Figure 1. Example annotations for a few seconds of audio. In this fictional clip, a female adult (FA1) speaks directly to the target child (CHI), a male adult (MA1) begins to speak to the female adult but cuts himself off, and a child is talking and singing, but it is unclear who his addressee is. All CHI utterances are annotated for vocal maturity (i.e. given canonical babble, check for use of recognizable words; given recognizable words, check for multi-word use). All non-CHI utterances are annotated for addressee type (C = child-directed speech / T can be used instead to indicate target-child directed speech specifically).